

McGINN & GIBB, PLLC
A PROFESSIONAL LIMITED LIABILITY COMPANY
PATENTS, TRADEMARKS, COPYRIGHTS, AND INTELLECTUAL PROPERTY LAW
8321 OLD COURTHOUSE ROAD, SUITE 200
VIENNA, VIRGINIA 22182-3817
TELEPHONE (703) 761-4100
FACSIMILE (703) 761-2375; (703) 761-2376

**APPLICATION
FOR
UNITED STATES
LETTERS PATENT**

APPLICANT: Tetsuo Shibuya et al.

FOR: **SYSTEM AND METHOD FOR
IDENTIFYING GENES**

DOCKET NO.: **YOR920010126US2**

SYSTEM AND METHOD FOR IDENTIFYING GENES

CROSS REFERENCE TO RELATED APPLICATIONS

This Application claims the benefit of Provisional Application No. 60/265,553 which was
5 filed on February 2, 2001 by Isidore Rigoutsos, et al., assigned to the present assignee, and which
is incorporated herein by reference.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to a system and method for identifying genes and,
10 more particularly, a system and method which utilizes a database of patterns to identify genes.

Description of the Related Art

Gene identification is one of the most important problems in molecular biology and has
15 been receiving increasing attention with the advent of automated large scale sequencing projects.
Indeed, more than 70 complete genomes currently exist in the public domain, while the
sequencing of many others is currently in progress. Consequently, the automated identification
of the protein coding regions in a newly sequenced genome is gaining importance.

Accurate gene prediction is of relevance to many biological applications. For instance,
20 the predicted coding regions can be used to generate probes for a DNA microarray, or to form the
basis for knockout experiments. In addition, the candidate proteins that correspond to these
predicted genes might be used as new drug targets, and so forth.

Specific attention has been given to the prokaryotic gene identification problem. With the exception of a handful of reported instances in archaeal organisms, splicing generally does not occur in prokaryotes and thus the problem of gene identification in these organisms is assumed to be simpler than its eukaryotic counterpart. Even so, the available schemes for the in silico gene prediction on prokaryotic genomes can be improved further and increasingly accurate prediction methods are always sought.

Over the years, a large number of methods have been proposed that address the gene identification problem. These methods can be largely divided into two categories. The first school of thought makes use of the statistics of DNA sequences to determine gene locations. It was observed early on that the nucleotide usage exhibits different statistical properties in DNA regions that code for genes than it does outside: the concept of the CpG island (e.g., see Bird, A., (1987) "CpG islands are gene markers in the vertebrate nucleus", *Trends in Genetics*, 3: 342-347) is a demonstration of such a difference in statistical behavior.

Among the gene identification methods that make use of this observation, hidden Markov models (HMMs) are probably the most popular. Specifically, HMMs are used in conventional methods such as GLIMMER (e.g., see Delcher, A. L., et al (1999), "Improved Microbial Gene identification with GLIMMER", *Nucl. Acid. Res.*, 27 (23): 4636-4641; and Salzberg, S. L., et al., (1998) "Microbial Gene Identification Using Interpolated Markov Models", *Nucl. Acid. Res.*, 26(2): 544-548) and GeneMark (Lukashin, A. V., and Borodovsky, M., (1998), "GeneMark.hmm: New Solutions for Gene Identification", *Nucl. Acid. Res.*, 16(4): 1107-1115).

The second school of thought advocates a strategy that is based on similarity searches in databases containing genomic information (e.g., see Badger, J. H. and Olsen, G. J., (1999),

“CRITICA: Coding Region Identification Tool Invoking Comparative Analysis”, *Molecular Biology and Evolution*, 16: 512-524; Bafna, V., and Huson, D. H., (2000), “The Conserved Exon Method for Gene Finding”, *Proc. ISMB ‘00*; Gelfand, M. S., Mironov, A. A., and Pevzner, P., (1996) “Gene Recognition Via Spliced Alignment”, *Proc. Natl. Acad. Sci. USA*, 93: 9061-9066; Gish, W., and States, D. J., (1993) “Identification of Protein Coding Regions by Database Similarity Search”, *Nat. Genet.*, 3: 266-272; and Robinson, K., Gilbert, W., and Church, G., (1994) “Large-scale Bacterial Gene Discovery by Similarity Search”, *Nat. Genet.*, 7: 205-214).

Here one searches in existing databases for either proteins or DNA regions in other genomes that share similarities with candidate proteins corresponding to open reading frames (ORFs) identified in the genome under consideration (e.g., see Burge, C., and Karlin, S., (1998), “Finding the Genes in Genomic DNA”, *Current Opinion in Structural Biology*, 8: 346-354; Burset, M. and Buigo, R., (1996) “Evaluation of Gene Structure Prediction Programs”, *Genomics*, 34: 353-367; Claverie, J. M., (1998), “Computational Methods for Exon Detection”, *Molecular Biotechnology*, 10: 27-48; Claverie, J. M., (1997), “Computational Methods for the Identification of Genes in Vertebrate Genomic Sequences”, *Human Molecular Genetics*, 6(10): 1735-1744; Fickett, J. W., (1996), “The Gene Identification Problem: An Overview for Developers”, *Computers Chem.*, 20(1): 103-118; and Fickett, J. W. and Hatzigeorgiou, A. G., (1997), “Eukaryotic Promoter Recognition”, *Genome Research*, 7: 871-878).

However, these conventional strategies have shortcomings. Statistical methods like HMMs can find regions whose statistical behavior is similar to that of the used training set. But if no appropriate training sets are available, one must resort to using training sets that are derived through database search, or simply assume very long open reading frames to be coding for genes.

The statistics of coding regions often differ from organism to organism, and ideally one ought to use HMMs whose parameters are organism-dependent if one wishes to achieve high prediction ratio. That is, one must train HMMs separately for each genome.

It has also been demonstrated that there exist many genes that are statistically distinct from other genes of the same organism, such as genes that are the result of horizontal transfer (e.g., see Kehoe, M. A. , Kapur, V., et al., (1996) "Horizontal Gene Transfer Among Group A Streptococci: Implications for Pathogenesis and Epidemiology", *Trends Microbiol.*, 4(11): 436-443; and Nielsen, K. M., bones, A. M., et al., (1998), "Horizontal Gene Transfer From Transgenic Plants to Terrestrial Bacteria - A Rare Event?", *FEMS Microbiol. Rev.*, 22(2): 79-103). Such cases typically pose challenges to statistical methods.

Finally, short genes (e.g. fewer than 60-80 a.a.) cannot be predicted easily using statistical methods. Similarity-based methods are more successful in finding short genes or genes that are statistically different from those in the rest of the organism under consideration as long as similar genes or proteins already appear in the databases being searched. Additional problems arise if the shared similarity between a candidate gene and its database counterpart is very low. On the flip-side, there is no dependence of the quality of answers on the choice of training sets.

Similarity-based methods generally have an improved ability in determining the correct location of genes over statistical methods, a desirable property. It is for these reasons that large genome sequencing projects often employ a combination of methods from both schools (e.g., see Fleishman, R. D., et al., (1995), "Whole-genome Random Sequencing and Assembly of *Haemophilus Influenzae*", *Science*, 269: 496-512).

SUMMARY OF THE INVENTION

It is, therefore, an object of the present invention to provide a structure and method for accurately and efficiently identifying genes in DNA sequences.

5 The present invention includes a system for identifying genes which includes a pattern database including patterns of amino acids, an input device for inputting a DNA sequence, and a processor for processing the DNA sequence and patterns to identify a putative gene. For instance, the processor may determine possible open reading frames (ORFs) in the DNA sequence, generate an amino acid translation for each ORF, and identify a match of a pattern in the amino acid translation. The inventive system may report an ORF as a putative gene when one or more pattern matches are identified in the amino acid translation. An ORF includes a portion of the DNA sequence between a start codon and a stop codon.

10 The patterns may be derived from a parent database of one or more proteins and/or protein fragments. Further, the patterns may be generated (e.g., by the processor) from the amino acid sequences of the proteins and protein fragments in the parent database using a predetermined algorithm, such as the Teiresias Algorithm.

15 The inventive system may further include a memory device for storing data and instructions to be executed by the processor, and a display device for displaying an output from the processor.

20 Further, each pattern may be assigned a weight (e.g., depending upon how relevant the pattern is in determining whether an ORF is a putative gene). For instance, the processor may assign a weight to a given pattern in the pattern database. In addition, an occurrence of a pattern match in said amino acid translation may be identified with the help of a predetermined

algorithm (e.g., a pattern matching algorithm) used to identify matches of a ~~the~~ pattern in the amino acid translation.

The present invention also includes an inventive method of identifying genes which includes optionally generating a set comprising patterns of amino acids, computing an open
5 reading frame (ORF) in a DNA sequence, generating an amino acid translation for each ORF; and identifying potential matches of the patterns in the amino acid translation. The employed collection of patterns may be generated, for example, using a predetermined algorithm, such as the Teiresias algorithm to process a parent database of one or more amino acid sequences or fragments. Further, an ORF may be reported as a putative gene when one or more pattern
10 matches are identified in said amino acid translation.

The inventive method may also include assigning a weight to each pattern depending upon how relevant the pattern is in determining whether an ORF (e.g., an ORF in which the pattern is identified) is a putative gene. The method may also include displaying or printing
instances of a pattern (e.g., pattern matches) in the amino acid translation..

15 The present invention also includes a programmable storage medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform the inventive method for identifying genes.

With its unique and novel features, the present invention provides a novel system and method which accurately and efficiently identifies genes. The present invention may be
20 considered to combine the best characteristics of statistical approaches and database similarity searches.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

Figure 1 illustrates a system 100 for identifying genes according to the present invention;

Figure 2 is a flowchart illustrating a method 200 of identifying genes according to the present invention;

Figure 3 illustrates an exemplary embodiment of the inventive method 200 according to the present invention;

Figure 4(a) provides Table 1 which describes the seventeen genomes studied in the inventors' experiments using the present invention;

Figure 4(b) provides Table 2 which displays results generated in experiments using the present invention;

Figure 5 illustrates a typical hardware configuration 500 which may be used for implementing the present invention; and

Figure 6 illustrates a signal bearing medium 600 for performing a method of identifying genes according to the present invention.

DETAILED DESCRIPTION OF PREFERRED

EMBODIMENTS OF THE INVENTION

Referring now to the drawings, Figure 1 illustrates an inventive system 100 for identifying genes according to the present invention.

As shown in Figure 1, the inventive system 100 includes a pattern database 110 which includes patterns of amino acids. The patterns may be provided to system 100 or, optionally, the patterns may be derived from a database(s) comprising one or more amino acid sequence or fragment of an amino acid sequence, or otherwise be made available to the inventive system 100.

5 The system 100 also includes an input device 120 for inputting data (e.g., a given DNA sequence) and instructions, and a processor 130 for processing the DNA sequence and patterns to identify a putative gene.

Specifically, the processor 130 may process input data to determine open reading frames (ORFs) in the DNA sequence, generate an amino acid translation for each ORF, and identify matches of the patterns in the amino acid translation. Optionally, the processor 130 may be used to also derive a database of patterns to be used with the present invention, by processing, for example, a database comprising one or more amino acid sequence or amino acid sequence fragment (e.g., proteins and/or protein fragments) with a pattern discovery algorithm such as the Teiresias algorithm.

15 The inventive system 100 provides a new approach for tackling the gene identification problem. The approach employs a pattern database 110 (e.g., a database of patterns that may or may not cover all of the currently available sample of natural protein sequence space) to determine gene candidates among the ORFs that can be identified in a given DNA strand. Further, the inventive system 100 combines the best characteristics from each of the above-mentioned schools of thought (e.g., statistical methods and similarity-based methods). In addition, the inventive system 100 may associate the patterns in the pattern database 110 with appropriately computed weights which leads to further improvements in the gene identification

ability.

The concept of the pattern database was introduced in a number of publications on the “IBM Bio-Dictionary(TM)” work including Rigoutsos, I., Floratos, I., et al., (1999) “Dictionary Building via Unsupervised Hierarchical Motif Discovery” *Journal of Proteins: Structure, Function and Genetics*, 37 (2) (hereinafter “Article 1”); Rigoutsos, I., Floratos, I., et al., (2000) “The Emergence of Pattern Discovery Techniques in Computational Biology”, *Journal of Metabolic Engineering*, 2(3), 159-177 (hereinafter “Article 2”); and Rigoutsos, I., Gao, Y., et al., (1999), “Building Dictionaries of 1D and 3D Motifs by Mining the Unaligned 1D Sequences of 17 Archaeal and Bacterial Genomes”, *Proc. ISMB '99* (hereinafter “Article 3”), which are all incorporated herein by reference.

The pattern database may be created by using, for example, the Teiresias algorithm, which is explained in Rigoutsos, I., and Floratos, A., (1998) “Combinatorial Pattern Discovery in Biological Sequences: The Teiresias Algorithm”, *Bioinformatics*, 14(1): 55-67 and Rigoutsos, I., and Floratos, A., (1998), “Motif Discovery Without Alignment or Enumeration”, *Proc. 2nd ACM International Conference on Computational Molecular Biology (RECOMB '98)*, which are incorporated herein by reference.

For instance, if Σ denotes the alphabet of all 20 amino acids, when processing an input dataset containing a collection of strings from Σ^+ with the Teiresias algorithm, one can succinctly capture the patterns that can be discovered with the regular expression $\Lambda(\Lambda \cup \{".\})^* \Lambda$ where $\Lambda = (\Sigma \cup \Sigma \Sigma^* \Sigma)$. In this expression, ‘.’ is a “don't care” character which can be replaced by any character in Σ . That is, the generated patterns can either be a single alphabet symbol, or strings that begin and end with a symbol or a bracket with two or more characters, and contain an

arbitrary combination of zero or more residues, brackets with at least two alphabet characters, and don't care characters. A bracket is meant to denote a "one of" choice. In other words, for example, [CPM] denotes exactly one of C, P or M. Also, a bracket can have a minimum of 2 (two) alphabet characters but obviously not more than $|\Sigma| - 1$.

5 A pattern t is called an $\langle L, W \rangle$ pattern (with $L \leq W$) if every substring of t of length W comprises L or more non-don't care positions. The smallest length of an $\langle L, W \rangle$ pattern is obviously equal to L whereas its maximum length is unbounded. Any given choice for the parameters L and W has a direct bearing on the degree of remaining similarity among the instances of the sequence fragments that the pattern captures. Thus, the smaller the value of the ratio L/W , the lower the degree of local similarity. Associated with each pattern t is its support which is denoted by K and represents the minimum number of instances of a pattern t in the input database from which it was derived.

The patterns of amino acids stored in the pattern database 110 are commonly referred to simply as patterns. As noted above, the patterns may completely describe and account for the currently known sequence space of natural proteins at the amino-acid level. However, this is not a necessity for this algorithm to operate. The patterns may be derived, for example, by processing a large public database (e.g., GenPept or SwissProt) of proteins and protein fragments using the Teiresias algorithm and discovering patterns that occur a certain number of times (e.g., discovering all $\langle 6, 15 \rangle$ patterns that occur 2 or more times, i.e. $L=6$, $W=15$, and $K=2$).

20 Alternatively, the patterns may be provided to the algorithm through other means, e.g., through access to an existing collection of patterns such as those contained in the PROSITE database. The availability of such a collection of patterns permits a user to effectively and successfully

tackle a number of tasks including, for example, similarity searching (e.g., see Floratos, A., Rigoutsos, I., et al., (1999) "Sequence Homology Detection Through Large Scale Pattern Discovery", *Proc. RECOMB '99*), functional annotation (e.g., see Article 3), phylogenetic domain analysis (e.g., see Article 2), as well as gene identification.

5 For example, in Article 1, the inventors described how to compute a pattern database from the GenPept release from February 10, 1999 that contains ~387,000 sequences with a grand total of ~120M amino acids. The computation gave rise to a pattern database 110 that comprised ~26M patterns and which accounted for (i.e. covered) 98.12% of the amino acid positions in the processed input.

10 As explained above, the pattern database 110 may substitute a given sequence database of proteins and fragments by a collection of patterns (i.e. regular expressions) that represents combinations (e.g., patterns) of amino acids that appear two or more times in the processed input.

15 Therefore, the inventive system 100 is able to successfully tackle the problems sought by researchers for several reasons. For instance, the pattern database 110 could be extracted from a large and diverse collection of proteins and protein fragments which are readily available, given the currently large number of completed and ongoing genome sequencing projects which contribute to the public databases (e.g., see Article 2). In other words, the invention may assume that a set of patterns has been made available somehow and simply uses these patterns.

20 Referring again to Figure 1, the inventive system also includes an input device 120. The input device 120 (e.g., a keyboard) may be used, for example, to input data (e.g., data generated by the user, or downloaded from another database such as a public database) to the inventive system 100, and for inputting instructions for processing the input dataset by the processor 130.

For example, the input device 120 may be used to input a DNA sequence.

The inventive system 100 may also include a memory device (e.g., RAM, ROM, etc.) which stores input data and instructions for processing such data. For example, such data may be downloaded from another database via the World Wide Web (e.g., Internet) into the memory device. For example, a DNA sequence to be studied for the presence of genes in it may be stored in the memory device (e.g., RAM, ROM, etc.) in the inventive system 100. Alternatively, data (e.g., a dataset) may be downloaded to the inventive system 100 directly from another database, for example, over the World Wide Web (e.g., Internet) and processed by the inventive system 100 without being permanently stored therein.

As shown in Figure 1, the inventive system 100 also includes a processor 130 (e.g., a microprocessor) which may be used to process the data which is input into the inventive system 100. The processor 130, may translate an input DNA sequence into an amino acid translation. Further, the processor 130 may perform a process (e.g., a pattern matching process) in order to match the patterns of amino acids stored in the pattern database 110 with the amino acid translation from the input DNA sequence.

In short, the processor 130 may process the input DNA sequence to be studied by determining all possible open reading frames (ORFs) in the DNA sequence. An ORF may include, for example, the DNA sequence between a start codon and a stop codon. For example, the processor 130 may compute all possible ORFs in each of the reading frames (e.g., three reading frames), and both for the forward and reverse strands of the given DNA sequence. The number of truly coding regions will be a proper subset of this collection of ORFs.

For each ORF, the processor 130 may generate an amino acid translation. If the ORF

under consideration is indeed a coding sequence, then the instances of one or more of the patterns from the pattern database 110 should be identifiable in the ORF's translation, and *vice versa*. If the number of patterns whose located instances exceeds a predetermined threshold, the ORF may be reported as a putative gene. Further, the higher the number of patterns that can be found in a given ORF and the more evenly their instances are distributed over the ORF's translation, the more likely it is that the ORF under consideration is a coding sequence.

To improve the efficiency of the inventive system 100 (e.g., to improve the gene-identifying capabilities of the inventive system 100) an optional weighting scheme may be included. In other words, in addition to the number of patterns that can be located within the translation of an ORF, the very nature of these patterns carries weight when deciding whether the ORF is indeed a coding one. In general, any two patterns that will match an amino acid translation will affect this decision differently. By summing up the scores of the patterns matching an ORF, a quality measure can be determined that will allow the ORF to be characterized as a putative gene, or otherwise.

The processor 130 may, therefore, optionally weight the patterns in the pattern database 110. For instance, where $T = \{t_1, t_2, \dots, t_n\}$ is the complete collection of patterns in the pattern database 110, if a putative protein s is coded for by some ORF from a given DNA sequence, and l is the length of s , it could be said that a pattern matches at position j of the amino acid sequence s if an instance of the pattern can be found beginning at the j -th location of s .

For example, G..G.GK[ST]TL matches the sequence MTHVLIKGAGGSGKSTLAFW beginning at position 8 of the sequence. Further, letting T_{sj} denote the set of patterns that match beginning at position j of s , letting T'_{sj} denote the set T/T_{sj} , letting $T_s = \{t_{v1}, t_{v2}, \dots, t_{vm}\}$ denote the

concatenated list of T_{s_j} 's for all j ($1 \leq j \leq l$), it can be seen that T_{s_j} 's for different j 's can contain the same pattern, thus T_s is in general a multiset.

The inventive system may, therefore, determine a coding quality measure for an ORF under consideration based, in part, on the probability, p_i , that pattern t_i matches an actual amino acid sequence at a fixed location, and the probability, q_i , that t_i matches the amino acid translation of a non-coding ORF at a fixed location.

For instance, letting $w_i = \log p_i - \log q_i$ where w_i is the weight associated with pattern t_i , and considering the sum of weights of the patterns matching anywhere in the translation s of an ORF as the measure W_s that is characteristic of the coding quality of the ORF under consideration, in many cases the following equation can be used to express the coding quality measure of an ORF:

$$W_s = \sum_{i=1}^m w_{v_i} = \sum_{i=1}^m (\log p_{v_i} - \log q_{v_i}) = \log R_s^3$$

where R_s^3 is an approximation of the relative likelihood that two candidate ORFs are coding.

Further, at times, a situation may be encountered where multiple start codons match the same stop codon so that the appropriate start/stop pair must be chosen. A straightforward solution may involve picking the start codon which will result in the highest value for the coding quality measure. However, selecting the start codon in such a way will not necessarily result in the longest ORF because patterns can also have negative associated weights.

On a related note, ATG is the most frequently used start codon but not the only one. It is thus conceivable that the different start codons be treated in a non-uniform manner. For example, if $\{c_1, c_2, \dots, c_k\}$ denotes the set of possible start codons, and f_i is the probability that c_i is

the start codon of a randomly chosen coding region, and if f'_i is the probability that c_i is observed in non-coding regions, and g_i is given as $\log f_i - \log f'_i$, the term $W_s + g$ may be used (instead of W_i) as the measure of coding quality for the amino acid translation s of an ORF that is initiated by the start codon c_i .

5 The inventive system 100 may, therefore, optionally utilize the values for p_i 's and q_i 's to compute the coding quality measure. One way to compute these p_i 's and q_i 's is to compute them with the help of known actual genes and non-coding ORFs. Further, the values of f_i 's and f'_i 's can also be obtained from the same training set.

10 But these values may also be obtained in the absence of a training set. For example, for p_i values, the probabilities computed with the help of the protein database from which the pattern database is derived can be used. Alternatively, the values can be computed using, for example, very long ORFs instead of actual coding regions. In addition, to obtain q_i values, non-ORF regions can be used, or the values may be obtained by estimating the probability of random occurrence based on an appropriately chosen amino acid bias.

15 When each ORF has been associated with its coding quality measure, the inventive system 100 may determine which ORFs correspond to putative genes by appropriately setting a threshold value. Such a threshold value may be input by a user using the input device 120 or may, for example be stored in a memory device accessible by the processor 130. The higher the value of the measure for a given ORF the more likely it is that it is a coding one.

20 Further, the inventive system 100 may use the processor 130 to identify matches of the patterns from the pattern database in the amino acid translation from an input DNA sequence by using, for example, a nested loop or other method. Alternatively, a predetermined algorithm may

be used to identify matches of the patterns in the amino acid translations.

For instance, using a predetermined algorithm, the processor 130 in the inventive system 100 may compare the patterns in the pattern database 110 with the amino acid translations of a large number of ORFs for each complete genome. This operation may be carried out as efficiently as possible so as to reduce computing time. The inventive system 100 may, therefore, include an algorithm (e.g., an algorithm stored in a memory device accessible by the processor 130) for performing such comparison.

Therefore, in summary, the processor 130 in the inventive system 100 may determine possible open reading frames (ORFs) in the DNA sequence, generate an amino acid translation for each ORF, and identify matches of the patterns in the amino acid translation. Further, the processor 130 may output the results of such computations to a database (e.g., a memory device such as RAM, ROM etc.). Further, the results may also be output to a display device (e.g, video display device) or printer for analysis by a user.

Referring again to the figures, Figure 2 is a flowchart illustrating an inventive method 200 for identifying genes according to the present invention. As shown in Figure 2, the inventive method 200 includes providing (210) a pattern database comprising patterns of amino acids. For instance, the pattern database may be derived from a database of one or more amino acid sequence or amino acid sequence fragment (e.g., proteins or protein fragments). The inventive method 200 also includes computing (220) all possible open reading frames (ORFs) in a DNA sequence, generating (230) an amino acid translation for each ORF, and identifying (240) matches of patterns from the pattern database in the ORF's amino acid translation.

The inventive method 200 may be more clearly understood by referring to Figure 3 which

illustrates an example of the inventive method 200 as it is used to identify genes in a given DNA sequence. As shown in Figure 3, the inventive method 200 provides (210) a pattern database (e.g., from a database (e.g., a public database) of amino acids or amino acid fragments (e.g., proteins and/or protein fragments)). The inventive method 200 also computes (220) possible open reading frames (ORFs) in a DNA sequence, and generates (230) an amino acid translation (e.g., a candidate gene) for each ORF.

The inventive method 200 may, thus, identify (240) matches of patterns from the pattern database in the amino acid translation, for example, by locating instances (e.g., matches) of patterns from the pattern database in the candidate gene and determining if support goes above the given threshold value. If yes, the ORF may be reported as a putative gene, and if not, the inventive method 100 proceeds with the next ORF.

Experiments and Results

Experiments conducted by the inventors have confirmed the efficacy and efficiency of the present invention. For example, the inventors have applied the present invention and gene-finding algorithm to several archaeal and bacterial genomes.

In the experiments, the inventors generated a pattern database with the help of the *Teiresias* algorithm. Specifically, an instance of the pattern database known as the IBM Bio-Dictionary was computed for the November 15, 1999 release of the GenPept database which contained 448,290 proteins and protein fragments corresponding to a total of 122,609,801 amino acids (this was the same process explained in Article 1). The pattern database generated by the inventors from the November 15, 1999 release of GenPept contained 31,184,670 patterns (e.g.,

small sequences commonly referred to as “seqlets”). These patterns accounted for 98.10% of the amino acids in the processed database.

It should be pointed out that the GenPept release used to build the pattern database chronologically preceded the releases of several of the test genomes that the inventors have processed and on which the inventors ran their experiments. This was an intentional choice by the inventors and was meant to demonstrate the present invention’s extrapolation capability.

In their experiments, the inventors used the present invention to process seventeen (17) complete genomes. Of these genomes, four were archaeal (*A. fulgidus*, *M. jannaschii*, *M. thermoautotrophicum*, *P. abyssi*) whereas the remaining thirteen were bacterial. Table 1 in Figure 4(a) shows the list of the genomes as well as their lengths in nucleotides, the numbers of all identifiable ORFs that are longer than 60 nucleotides, and the numbers of annotated coding regions from each genome that have been included in the public databases. In Figure 4(a), it can be seen that the number of coding regions is roughly 1/1000-th of the length of the genome. It is also very likely that these genomes contain coding regions that have not yet been reported. Also, it must be remembered that these annotated coding regions are in reality *putative* and have been annotated by scientists typically in the absence of wet laboratory experiments.

The inventors first carried out experiments using pattern weights that were computed separately from the reported coding regions of each genome as listed in the public databases.

Table 2 in Figure 4(b) shows the results of the experiments using the present invention (e.g., using the inventive gene finding algorithm). Here, the open reading frames that occupy the top (a) 1.0 x #CDS and (b) 1.1 x #CDS positions when sorted in decreasing value of coding quality as potential coding regions are reported. In Figure 4(b), #CDS is the number of annotated

coding regions in the databases (see last column of Table 1 in Figure 4(a)). It should be noted that the number of not-previously-predicted annotated genes is equal to that of reported additional putative genes in the case “a” as shown in the left hand column of the table.

The 'Annotated' column in Table 2 (Figure 4(b)) shows the numbers of correctly predicted ORFs. In this case, ORFs which were reported as putative genes overlap with regions that have been designated as coding in the public databases. This result is also shown as a percentage of the genes that have already been reported in the genomic database entry for the respective genome (see, for example, the last column of Table 1 in Figure 4(a)).

The 'Additional' column shows the number of ORFs that the present invention reported as putative genes but for which there was no database entry characterizing them as such. The 'Hit' column shows how many of these “additional putative genes” have substantial similarity with proteins contained in the January 15, 2001 release of SwissProt/TrEMBL, which are found by FASTA with default options and reported to have E(.) values that are not larger than 1.0e-5. The 'Score' column shows the lowest value of the coding quality measure for the ORFs which were reported as putative genes.

As shown in Table 2 (Figure 4(b)), the present invention can achieve very high prediction rates. In particular, when the top 1.1 x #CDS positions are considered (e.g., case “b” above as shown in the right hand column of Table 2), the prediction ratios, (*i.e.* the rates of predicted ORFs among annotated coding regions) exceed the 94% mark for all of the genomes examined by the inventors. As a matter of fact, and with the exception of the *E. coli* and *Synechocystis sp.* genomes, the inventors' prediction rates exceed the 98% mark. For many of these genomes, additional putative genes (listed in the 'Additional' column) that almost invariably share

similarities with proteins from SwissProt/TrEMBL according to FASTA, are reported. These additional putative genes are likely to be coding with high probability. The coding quality measures that the present invention attributed to these ORFs is high enough to warrant laboratory experiments which can verify that the ORFs are indeed coding. It is notable that the present invention here achieved *perfect* prediction in the case of *B. burgdorferi*. Given that the inventors made no use of information regarding promoter regions, terminators or enhancers, the results achieved by the present invention were very encouraging.

It is worth stressing the fact that these high prediction rates were achieved by using the same universal set of pattern weights on all genomes. This is in marked contrast with methods that are based on statistical techniques, such as HMMs, where the user computes and uses genome-specific parameters, and is indicative of the potential of the present invention.

In summary, the present invention provides a new system and method for solving the gene identification problem. The invention relies in part on a straightforward idea and as the reported experimental results demonstrate, it can predict genes very accurately. So as to further demonstrate the capabilities of the invention, the inventors intentionally relied upon a pattern database that was built from a November 1999 release of GenPept and applied it to genomes whose ORFs were only in part included in GenPept or not included at all. It is easy to see how repeating the gene discovery process with a pattern database that is computed from a more recent release of a public database such as GenePept would further improve the quality of the inventors' experimental results.

In addition, the inventors could potentially augment the present invention by associating each of the patterns in the pattern database with manually-derived or automatically-derived

weights.

Further, it should be noted that in addition to correctly finding those of the ORFs that have already been reported in the public databases as putative genes, the present invention determines additional candidate genes in almost every single one of the genomes on which the inventors ran experiments. The inventors used FASTA to determine similarities between the additional candidate genes and the entries currently in SwissProt/TrEMBL and, in fact, such similarities were identified for many of these genes (see Table 2 in Figure 4(b)).

Referring now to Figure 5, system 500 illustrates a typical hardware configuration which may be used for implementing the inventive system and method for identifying genes. The configuration has preferably at least one processor or central processing unit (CPU) 511. The CPUs 511 are interconnected via a system bus 512 to a random access memory (RAM) 514, read-only memory (ROM) 516, input/output (I/O) adapter 518 (for connecting peripheral devices such as disk units 521 and tape drives 540 to the bus 512), user interface adapter 522 (for connecting a keyboard 524, mouse 526, speaker 528, microphone 532, and/or other user interface device to the bus 512), a communication adapter 534 for connecting an information handling system to a data processing network, the Internet, and Intranet, a personal area network (PAN), etc., and a display adapter 536 for connecting the bus 512 to a display device 538 and/or printer 539. Further, an automated reader/scanner 541 may be included. Such readers/scanners are commercially available from many sources.

In addition to the system described above, a different aspect of the invention includes a computer-implemented method for performing the above method. As an example, this method may be implemented in the particular environment discussed above.

Such a method may be implemented, for example, by operating a computer, as embodied by a digital data processing apparatus, to execute a sequence of machine-readable instructions. These instructions may reside in various types of signal-bearing media.

Thus, this aspect of the present invention is directed to a programmed product, including signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital data processor to perform the inventive method.

Such a method may be implemented, for example, by operating the CPU 511 to execute a sequence of machine-readable instructions. These instructions may reside in various types of signal bearing media.

Thus, this aspect of the present invention is directed to a programmed product, comprising signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital data processor incorporating the CPU 511 and hardware above, to perform the method of the invention.

This signal-bearing media may include, for example, a RAM contained within the CPU 511, as represented by the fast-access storage for example. Alternatively, the instructions may be contained in another signal-bearing media, such as a magnetic data storage diskette 600 (Figure 6), directly or indirectly accessible by the CPU 511.

Whether contained in the computer server/CPU 511, or elsewhere, the instructions may be stored on a variety of machine-readable data storage media, such as DASD storage (e.g., a conventional "hard drive" or a RAID array), magnetic tape, electronic read-only memory (e.g., ROM, EPROM, or EEPROM), an optical storage device (e.g., CD-ROM, WORM, DVD, digital optical tape, etc.), paper "punch" cards, or other suitable signal-bearing media including

transmission media such as digital and analog and communication links and wireless. In an illustrative embodiment of the invention, the machine-readable instructions may comprise software object code, compiled from a language such as "C," etc.

With its unique and novel features, the present invention provides a novel system and method which accurately and efficiently identifies genes. The present invention may be considered to combine the best characteristics of statistical approaches and database similarity searches.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims. For example, instead of a database of patterns of amino acids, alternatively a user could employ a database of patterns based on nucleotides (e.g., generated or provided) and match the patterns of the database in an open reading frame (ORF) so as to eliminate the need to translate the candidate ORF.